The Induced Bounded-Degree Subgraph Problem and Stream Control in MIMO Networks

Harish Sethu

Email: sethu@drexel.edu

Department of Electrical and Computer Engineering

Drexel University, Philadelphia, PA 19104-2875

Abstract

In this report, we consider maximal solutions to the induced bounded-degree subgraph problem and relate it to issues concerning stream control in multiple-input multiple-output (MIMO) networks. We present a new distributed algorithm that completes in logarithmic time with high probability and is guaranteed to complete in linear time. We conclude the report with simulation results that address the effectiveness of stream control and the relative impact of receiver overloading and flexible interference suppression.

I. INTRODUCTION

Multi-hop wireless sensor networks and mobile ad hoc networks have been suggested for a variety of applications in both military and civilian environments. In such networks, the aggregate traffic-carrying capacity is determined by the number of concurrent transmissions that can be supported without interference from other transmissions. If we consider each transmission link in a network as a vertex in a graph with an edge connecting each pair of vertices corresponding to interfering links, the problem of finding the maximum aggregate capacity reduces to the problem of finding the largest independent set in the graph [1], [2]. In this report, we pose a generalized version of the independent set problem inspired by the technology of multiple-input multiple-output (MIMO) enabled by smart antennas [3].

In a MIMO link, multiple antenna elements at both ends of a communication link allow simultaneous transmission of multiple data streams, thus achieving a higher capacity without increased bandwidth or power requirements. With q degrees of freedom (number of antenna elements), a MIMO receiver can isolate and decipher up to q streams that reach the receiver. A MIMO link, however, may choose only the strongest streams (channel modes) and transmit on fewer than q streams. Such *stream control*, illustrated in Fig. 1, is known to allow multiple interfering links operating simultaneously to achieve a better overall throughput than if each link transmits on all available streams [4], [5]. As identified in [6], there are at least two aspects related to stream control that influence, in opposite directions, the potential capacity of a MIMO network. The phenomenon of *receiver overloading* offsets the gains from stream control because an active receiver in a contention zone reduces the number of potential transmitters in all the contention zones that it belongs to. This can reduce the ability to exploit spatial diversity since more links are active when stream control is employed and therefore, more active receivers may be overloaded. On the other hand, stream control allows *flexible interference suppression* because the number of degrees of freedom required to suppress interference at a node may be smaller than the number of interfering streams. For example, it may be possible to decode two desired streams using up two degrees of freedom while suppressing two other undesirable interfering streams using up only one additional degree of freedom, thus using a total of only three (instead of four) degrees of freedom.

In this report, we seek a graph-theoretic approach to gain some insight into the relative influence of receiver overloading and flexible interference suppression on the effectiveness of stream control in a MIMO network. We relate the issue to the problem of finding a maximal induced bounded-degree subgraph of the contention graph of a MIMO network.

II. PROBLEM STATEMENT

Consider a MIMO network in which each node possesses q antenna elements (therefore, q degrees of freedom) and is capable of transmitting on q different streams to a neighboring node. Consider a graph $G = (V, E)$ in which each vertex represents a potential stream in a MIMO link. Two streams that interfere with each other are connected by an edge in the graph. Thus, the streams comprising a single MIMO link make up a q-vertex clique. Fig. 2(a) shows a simple MIMO network with five physical nodes. Assuming two streams per MIMO link, Fig. 2(b) shows the corresponding graph representing the interference pattern between the different streams. Our goal is to maximize the number of active streams while ensuring that there are no more than a certain number, say k, of streams interfering with any given active stream. For example, in

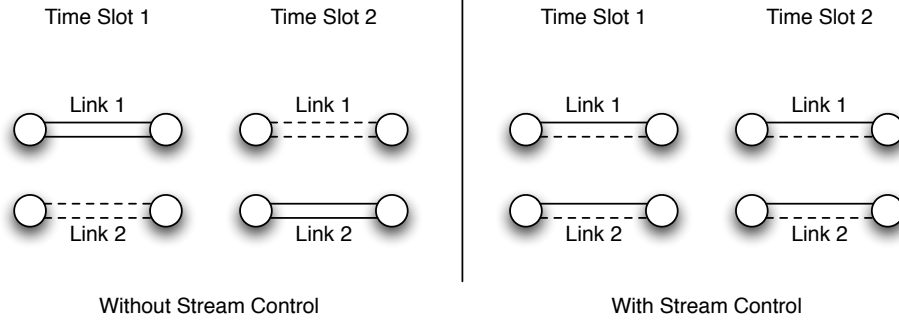| Time Slot 1 | Time Slot 2 | Time Slot 1 | Time Slot 2 |

Fig. 1. Two interfering MIMO links, each comprised of two streams, can achieve a higher data rate if both links transmit simultaneously on their stronger streams instead of taking turns and transmitting during different slots on all streams. Solid line indicates an active stream while a dashed line indicates an unused stream.

the case in which it takes exactly one degree of freedom at an active receiver to suppress an undesired interfering stream (i.e., no flexible interference suppression), an active stream should have no more than $k = q - 1$ other streams interfering at the intended receiver. The problem now becomes that of finding the largest subset of vertices $V' \subset V$ such that for any vertex $v \in V'$, there exist no more than k other vertices in V' connected to v by an edge in G. Thus, this becomes the maximum induced bounded-degree-k subgraph (IBDS(k)) problem which is known to be NP-complete for general graphs [7], although linear time solutions have been discovered for certain types of graphs [8]. The well-known maximum independent set problem is an instance of maximum IBDS(k) for $k = 0$. We hasten to add that maximum IBDS(k) is different from another better-known and well-studied generalization of the maximum independent set problem, the k-independent set problem (related to the k-coloring of graphs) [9].

We define a *maximal* IBDS(k) as an induced bounded-degree-k subgraph to which no additional vertex can be added. We also define two simple variations of the maximal IBDS(k) problem for relevance in practical MIMO networks. Streams corresponding to the same MIMO link (i.e., between the same pair of nodes) are said to belong to the same *family*. Two potential streams with at least one physical node in common are said to belong to the same *superfamily*. In MIMO networks, practical considerations often impose that a node does not use two streams corresponding to different MIMO links simultaneously. We define the maximal IBDS-R(k) problem as the maximal IBDS(k) problem with the additional restriction that for any two vertices $v_1, v_2 \in V'$, if v_1 and v_2 belong to the same superfamily, then they belong to the same family. A second variation is motivated by the fact that exploiting flexible interference suppression requires that each MIMO link operate only a subset (say, no more than g) of the maximum number of streams possible and employ the remaining available degrees of freedom for interference suppression. We define the maximal IBDS-R(k, g) problem as the maximal IBDS-R(k) problem with the restriction that no more than g vertices from the same family are selected into the maximal subgraph.

III. The Distributed Algorithm

Fig. 3 presents a pseudo-code description of one round of DML-IBDS(k), a new distributed randomized algorithm for the maximal IBDS(k) problem (DML stands for distributed maximal).
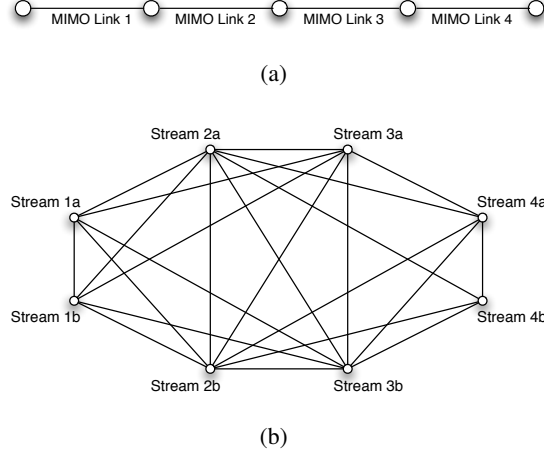
3

(a)

(b)

Fig. 2. Example illustrating (a) A simple network of MIMO nodes, and (b) the corresponding graph on which the distributed induced bounded-degree subgraph algorithm is run.

Each vertex v holds four variables, $r(v)$, $b(v)$, $\hat{b}(v)$ and $a(v)$. In the beginning before the first round, each vertex v computes a unique random integer $r(v)$. As in most distributed algorithms for networks, we assume that the vertices possess unique ids and therefore, the uniqueness of the random numbers is readily ensured by appending the vertex id to the generated pseudo-random number. $r(v)$ is updated in subsequent rounds and is not necessarily a unique number except at the beginning of the first round. $b(v)$ is initially set to 0 and indicates whether or not vertex v, in some round, has $r(v)$ smaller than $r(u)$ for all neighbors u (this is what qualifies vertex v for inclusion in the subgraph). $\hat{b}$ is initially set to -1. When $\hat{b}$ is 1, it indicates that the vertex is chosen for inclusion in the subgraph and when it is 0, it indicates that the vertex has been eliminated from consideration. $a(v)$ is initially set to $k+1$ and represents the maximum number of additional vertices that may be chosen for the subgraph out of the set comprising of v and its neighbors.

Each vertex v sends its random number, $r(v)$, to its neighbors and receives numbers from its neighbors (lines 08–09). If none of v's neighbors carries a smaller value of r, then vertex v sets its $b(v)$ to 1 (lines 10–11). Neighbors communicate their value of b to each other (lines 12–13). A vertex that has its b value set to 1 will now consider updating its r value. If v has no neighbors with the same r value as itself or if all neighbors of v with the same r value are already chosen for inclusion into the subgraph (i.e., $b = 1$), then v updates its r value to the smallest among the received values but larger than $r(v)$ (lines 15–16). This step ensures that no more than one of the neighbors of a vertex with $b = 1$ will be selected into the subgraph in the same round (thus guaranteeing that $a(v)$ does not suddenly reach below zero; a vertex v with $b(v) = 1$ that does not yet have $a(v) = 0$ will always possess an $r(v)$ equal to the smallest of its neighbors, and therefore, only its smallest neighbor will potentially be chosen into the subgraph.)

Theorem 1: The DML-IBDS(k) algorithm finds a maximal subgraph with degree bound k in $O(n)$ rounds where n is the number of vertices in the graph.

Proof: At the beginning of a round, consider the vertex v with the smallest value of $r(v)$ such that $\hat{b}(v) = -1$. By lines 15–16, any vertex u with $\hat{b}(u) = 1$ that is still participating in

4

```
01: Initialization (at vertex v):
02:     r(v) = pseudo-random integer
03:     b(v) = 0
04:     b̂(v) = −1
05:     a(v) = k + 1

06: One round of DML-IBDS(k) (at vertex v):
07:     if a(v) > 0:
08:         Send r(v) to neighbors
09:         Receive r(u) from each neighbor u
10:         if r(v) ≤ r(u) for every neighbor u:
11:             b(v) = 1
12:         Send b(v) to neighbors
13:         Receive b(u) from each neighbor u
14:         if b(v) = 1:
15:             if ∄u such that b(u) = 0 and r(u) = r(v):
16:                 r(v) = min{r(u)|r(u) > r(v)}
17:         if b̂(v) = −1 and b(v) = 1:
18:             b̂(v) = 1
19:             if b(u) = 1 for every neighbor u:
20:                 Halt
21:         Decrement a(v)
22:         Decrement a(u) at each neighbor u
23:     if a(v) ≤ 0:
24:         if b̂(v) = 1:
25:             for each neighbor u:
26:                 a(u) = 0
27:                 if b̂(u) = −1:
28:                     b̂(u) = 0
29:         Halt
```

Fig. 3. Pseudo-code description of DML-IBDS(k). A vertex v is selected into the subgraph if $\hat{b} = 1$.

the rounds will have $r(u)$ no smaller than $r(v)$. Therefore, vertex v will have $\hat{b}$ set to 1 by the end of the round. Thus, in each round, at least one additional vertex has its $\hat{b}$ set to one, and therefore, the algorithm completes in $O(n)$ rounds.

A vertex v reduces $a(v)$ by one whenever itself or one of its neighbors is chosen into the subgraph (lines 21–22). However, $a(v)$ is never allowed to go below zero for vertices that have $\hat{b}$ set to 1 (i.e., vertices corresponding to streams chosen for the subgraph). Since a vertex halts when $a(v) = 0$ (line 29), the induced subgraph will satisfy the desired degree bound. Also, it is a maximal subgraph because a vertex halts only when $a(v) \leq 0$ (line 29) or when the vertex along with all its neighbors has been chosen for inclusion into the subgraph (line 20). ∎

In real networks, after a topology control algorithm such as RNG has been executed, each physical node would have a bounded number of MIMO links. The corresponding graph of
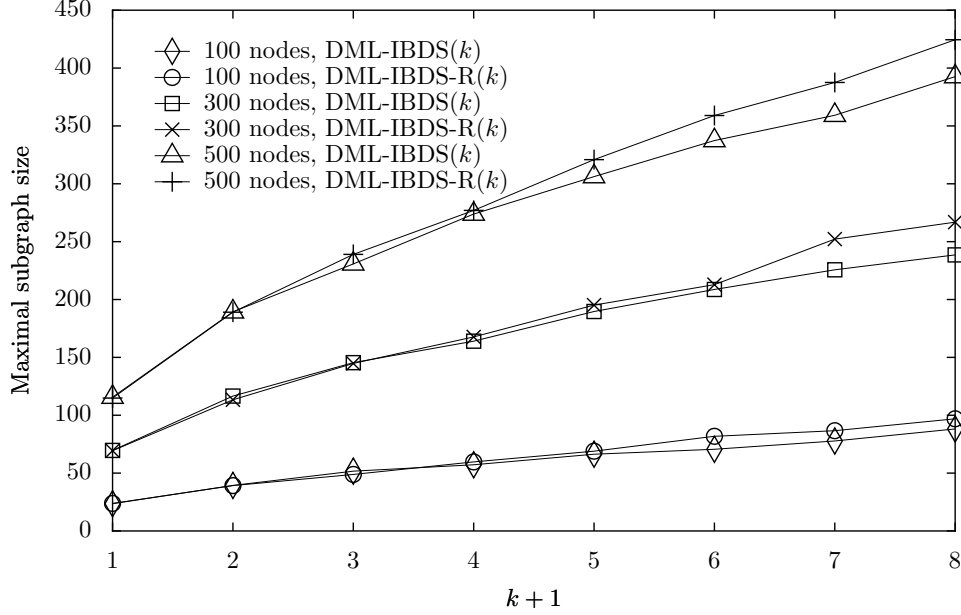
Fig. 4. Size of the subgraph with DML-IBDS(k) and DML-IBDS-R(k) plotted against $k + 1$. Note that in the absence of flexible interference suppression, $k + 1$ may be thought of as the number of antenna elements per node.

interference pattern among the streams, therefore, has a degree bound as well. The probability that any given vertex v ends up with an $r(v)$ no smaller than that of any of its neighbors is less than or equal to $1/(1 + deg_G(v))$. Given a bound on the degree, each node has a probability greater than some constant of having its $a(v)$ reduced by one. Since each of the n vertices begins with $a = k$ degrees of freedom, even though the worst-case number of rounds is $O(n)$, the actual number of rounds is $O(\log nk)$ with high probability (i.e., with probability greater than or equal to $1 - 1/(nk)^c$ for any $n \geq n_0$, $k > k_0$ for some positive constants c, k_0 and n_0).

Distributed solutions to variations of maximal IBDS(k) described in the previous section are only trivially different and are denoted in this report by DML-IBDS-R(k) and DML-IBDS-R(k, g).

IV. RESULTS AND CONCLUDING REMARKS

Fig. 4 presents simulation results on the performance of DML-IBDS(k) and DML-IBDS-R(k) as measured by the number of vertices in the degree-bounded subgraph that is generated. The networks used in our experiments consist of nodes located randomly in a unit square area and with the IBDS algorithms applied after the application of the DRNG (Directed Relative Neighborhood Graph) topology control algorithm [10]. We choose the DRNG protocol because it is one that can be employed in a multipath environment where a node cannot readily deduce the distance and direction of its neighbors. Each data point in this report represents an average of over 25 different randomly generated networks.

As can be seen from Fig. 4, DML-IBDS-R(k) performs slightly better than that for the unrestricted version, DML-IBDS(k), indicating that the negative impact of receiver overloading is reduced when the active streams at a node all belong to the same MIMO link. This suggests
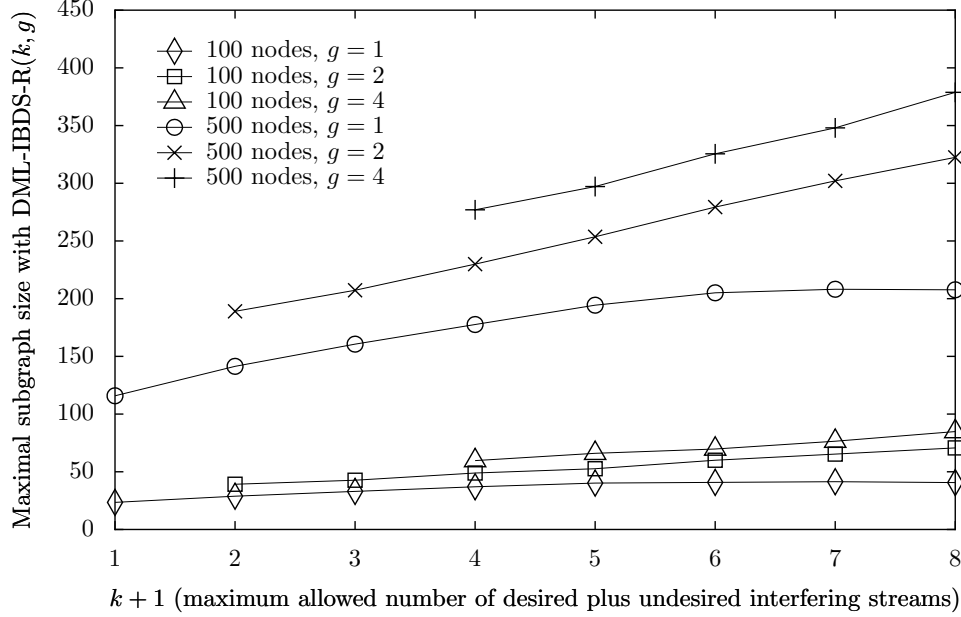
Fig. 5. Size of subgraph produced by DML-IBDS-R(k, g), i.e., when each MIMO link is allowed no more than g active streams and no more than $k + 1$ total streams (desirable plus undesirable) interfere at any receiver.

that the synchronization necessary to allow simultaneous use of streams to different neighbors at a node may not yield a higher aggregate throughput unless the stream control gains are high (i.e., if the strongest channel modes are significantly stronger than the weaker channel modes).

Fig. 4 also shows that the number of streams used (size of induced subgraph) does not increase linearly as the number of antenna elements increases. This is obviously sub-optimal to a trivial solution that involves no stream control: generate a maximal independent set, such as by running DML-IBDS(0), and at all the MIMO links chosen to transmit on a stream just allow all the streams to transmit. Since practical MAC algorithms do not usually approach the capacity corresponding to a maximal independent set, this sublinear increase in the number of active streams does not necessarily suggest that stream control is not desirable in practice. In fact, as shown in [6], stream control does improve performance in real contexts. Our results, however, shed light on the significant negative effect of receiver overloading and help us quantify the desired stream gains needed to offset the effect. For example, in a 500-node network in the absence of flexible interference suppression, the number of streams with one antenna element per node is 115 (subgraph size yielded by DML-IBDS-R(0)) while that with two antenna elements is 189 (subgraph size yielded by DML-IBDS-R(1)). This implies that for stream control to increase capacity in a network with two antenna elements per node, in the absence of flexible interference suppression, the average gain due to stream control should be about 22% ($2 \times 115/189 - 1$).

Fig. 5 shows us the gains achievable through flexible interference suppression and permit a quantitative comparison to the negative impact of receiver overloading. The subgraph size with DML-IBDS-R(k, g) indicates the number of active streams in the network when the number of active streams per MIMO link is bounded by g and the total number of streams (both desirable and the undesirable interfering ones) arriving at a receiver is bounded by $k + 1$. For example, for

7

500-node networks, Figs. 4 and 5 show that the subgraph size obtained with DML-IBDS-R(1) lies somewhere between those obtained by DML-IBDS-R$(4, 1)$ and DML-IBDS-R$(5, 1)$. If we have two antenna elements per node and if only one stream ($g = 1$) is used per MIMO link in order to exploit flexible interference suppression, the remaining one available degree of freedom has to be sufficient to suppress up to five additional undesirable streams to exceed the subgraph size that can be obtained by allowing both streams to operate whenever possible. These results suggest that, unless each available degree of freedom can be used to suppress several undesirable interfering streams, it is better to use any available degrees of freedom to active streams instead of flexible interference suppression. These insights offered by our algorithmic results can be used to design better MAC protocols for MIMO networks.

Acknowledgments

REFERENCES

[1] D. Peleg, *Distributed Computing: A Locality-Sensitive Approach*. SIAM, 2000.
[2] H. Balakrishnan, C. L. Barrett, V. S. A. Kumar, M. V. Marathe, and S. Thite, "The distance-2 matching problem and its relationship to the MAC-layer capacity of ad hoc wireless network," *IEEE Journal on Selected Areas in Communications*, vol. 22, no. 6, pp. 1069–1079, August 2004.
[3] S. Haykin and M. Moher, *Modern Wireless Communications*. Prentice Hall, 2005.
[4] M. F. Demirkol and M. A. Ingram, "Stream control in networks with interfering MIMO links," in *Proceedings of the IEEE Wireless Communications and Networking Conference*, 2003, pp. 343–348.
[5] R. S. Blum, J. H. Winters, and N. R. Sollenberger, "On the capacity of cellular systems with MIMO," *IEEE Communication Letters*, vol. 6, no. 6, pp. 242–244, June 2002.
[6] K. Sundaresan, R. Sivakumar, M. A. Ingram, and T.-Y. Chang, "Medium access control in ad hoc networks with MIMO links: Optimization considerations and algorithms," *IEEE Transactions on Mobile Computing*, vol. 3, no. 4, pp. 350–365, October-December 2004.
[7] J. A. Telle, "Complexity of domination-type problems in graphs," *Nordic Journal of Computing*, vol. 1, no. 1, pp. 157–171, Spring 1994.
[8] H. L. Bodlaender and B. de Fluiter, "Reduction algorithms for constructing solutions in graphs with small treewidth," in *Computing and Combinatorics*, ser. Lecture Notes in Computer Science, vol. 1090. Springer, 1996, pp. 199–208.
[9] D. B. West, *Introduction to Graph Theory*, 2nd ed. Prentice Hall, 2001.
[10] N. Li and J. C. Hou, "Localized topology control algorithms for heterogeneous wireless networks," *IEEE/ACM Transactions on Networking*, vol. 13, no. 6, pp. 1313–1324, December 2005.